



Institute for Scientific Computing Research

University
Collaborative
Research Program
Subcontract
Abstracts



**Center for Applied
Scientific Computing**

Structured Adaptive Mesh Refinement Running on Multi-Tier Computers

Scott Baden

University of California, San Diego

Abstract

We are investigating scalable parallel programming methodology and performance tradeoffs arising in structured adaptive mesh refinement running on multi-tier computers. Multi-tier computers, such as the ASCI Blue-Pacific machine, employ symmetric multiprocessor (SMP) nodes. Their performance tradeoffs are more of a challenge to master than previous-generation single-tier computers, which are based on single-processor nodes.

The program leverages previous work on the KeLP system, a C++ framework tailored to multi-tier architectures. We have experience with a "communication aware" load-balancing strategy that includes communication costs as part of the workload. We are refining the load-balancing strategy developed earlier, and exploring locality-enhancing and latency-hiding techniques on the ASCI Blue-Pacific platform. The investigation will deliver a computational testbed permitting LLNL scientists to explore portable scalable implementations of adaptive mesh applications running on a variety of platforms of interest to the Laboratory. The principal investigator will collaborate with John May of CASC to carry out performance analysis of the software infrastructure and to establish contact with potential Laboratory users.

Computational Fluid Dynamic Studies of Arterial Flow Disturbance Induced by Intravascular Stents

Abdul Barakat

University of California, Davis

Abstract

Atherosclerosis is an arterial disease whose pathological complications, namely heart disease and stroke, are the leading causes of mortality in the industrialized world. In its advanced form, atherosclerosis leads to plaques that protrude into arterial lumens and form stenoses, or even complete vessel occlusions that obstruct blood flow and give rise to the pathological events. One common interventional procedure involves the placement of an intravascular stent, an expandable wire mesh structure. The stent is introduced into the diseased artery in a compressed state and inflated at the stenosis or occlusion site to both restore blood flow and provide structural stability to the arterial wall. The major limitation to the success of this procedure, however, is restenosis, a complex and incompletely understood process by which plaques re-protrude into the vessel lumen within a period of a few months.

The placement of a stent in an artery mechanically damages the endothelium, the monolayer of cells lining the inner surface of all blood vessels. In vitro data indicate that the rate of endothelial repair after injury may be significantly slower in regions in which endothelial cells are exposed to relatively large fluid mechanical shear stress gradients, as occurs at the end points of flow separation zones. Therefore, flow separation in the vicinity of a stent may contribute to restenosis. The broad hypothesis driving our research is that the occurrence of flow separation depends on appropriate hemodynamic matching

between the stent design and the flow and geometric properties of the arterial segment in which the stent is positioned. We have been testing this hypothesis by using computational fluid dynamic techniques to probe the impact of various geometric and flow parameters on the occurrence of near-stent flow separation. Our results to date indicate that relatively small variations in stent wire thickness may have profound implications on the flow field within an artery and that thicker-wire stents are more likely to induce regions of flow separation with accompanying high wall shear stress gradients. This may partially explain the observation in animal studies of higher incidence of restenosis when using thicker-wire stents. We are extending the simulations to incorporate additional features that will render our numerical model more physiologically realistic. Our specific aims are (1) to perform three-dimensional steady and pulsatile flow simulations of the flow field in the vicinity of an intravascular stent positioned within a rigid-wall arterial segment with out-of-plane curvature, (2) to incorporate non-Newtonian fluid properties into the simulations, and (3) to incorporate certain aspects of arterial wall motion into the simulations.

The proposed collaboration with LLNL is critical for two reasons: LLNL's supercomputer are needed; and hydrodynamic simulation capabilities currently under development at LLNL may, in the future, allow the incorporation of additional physiological considerations into the aortic model, including fluid-wall coupling.

Implementing the Basic Self-Organizing Map (SOM) Algorithm for Data Mining

Jackson Beatty

University of California, Los Angeles

Abstract

We are developing a parallelized implementation of the basic self-organizing map (SOM) algorithm for data mining applications, designing a visually more effective SOM interface, and exploring the usefulness of volume SOM reference vector arrays. This research is important because SOMs have recently proven to be extraordinarily valuable tools for data mining in fields as diverse as mapping the distribution of atomic and molecular ions (Wolkenstein et al., 1997), identifying neuroactive compounds (Bucknecht et al, 1996), protein sequencing (Hanke et al., 1996) and galactic morphology (Naim et al., 1997).

We are implementing SOM procedures to study the problem of functional cortical organization by studying regularities in patterns of cerebral activation as evidenced in positron emission tomography and functional magnetic resonance imaging data. We believe that high-resolution self-organizing maps are best suited for our application, but high resolution can be produced only by using very large reference vector arrays. These are impractical using conventional serial computers, but are ideally suited for parallelization. We believe that parallelized SOM analysis, with improved visualization of the resulting SOM configurations, will not only materially aid us in our study of the functional cortical circuitry of the human brain, but will also contribute substantially to the computational infrastructure in a number of other scientific disciplines.

Parallel PIC Modeling of Semiclassical Quantum Models

Viktor Decyk and John M. Dawson

University of California, Los Angeles

Abstract

We are modeling many-particle quantum systems by combining a semiclassical approximation of Feynman path integrals with parallel computing techniques previously developed at UCLA for simulating plasmas.

Parallelized 3D Relativistic PIC Code for the Production of Useful Electron Bunches

Warren Mori

University of California, Los Angeles

Abstract

We are developing a parallelized three-dimensional relativistic particle-in-cell (PIC) code for studying the production of useful electron bunches using ultra-intense lasers-pulses. The electron bunches are produced when the radiation pressure of an intense laser either directly accelerates electrons or excites a plasma wave wake, which accelerates electrons. Ultra-short electron bunches could be useful for injection into high-energy particle accelerators, for radiation sources, and for the fast ignitor fusion concept. A reliable and robust code will be invaluable towards the realization of such ultra-short electron bunches in a laboratory.

We have carried out large physics simulations on the T3E at NERSC using an already existing code, PEGASUS, and have successfully developed a Fortran90-based object-oriented production code using the algorithms from PEGASUS. We will study the performance of a 3D parallelized version of this code and use it to investigate the production of useful electron bunches.

Algorithms and Software for Sensitivity Analysis of Large-Scale Differential–Algebraic Systems

Linda Petzold

University of California, Santa Barbara

Abstract

Sensitivity analysis of large-scale differential–algebraic (DAE) systems is important in many engineering and scientific applications. Sensitivity analysis generates essential information for design optimization, parameter estimation, optimal control, model reduction, management of uncertainty, process sensitivity and experimental design.

We are investigating and developing sensitivity analysis methods for large-scale DAEs. Our goal is the development of robust and efficient sensitivity software, based on the most recent versions of our widely employed codes DASSL and DASPK, and incorporating automatic differentiation software for evaluation of sensitivity derivatives. The proposed work is being done in close collaboration with Peter Brown and Alan Hindmarsh of LLNL, and includes both development and analysis of numerical methods and development of software.

Seismic Scattering for Strong Ground Motion Applications

Fred Pollitz

University of California, Davis

Abstract

Prediction of strong ground motion using forward computation of the propagating seismic wavefield is one of the few means by which to estimate the potential hazards posed by a future large earthquake. The most damaging ground motions typically have a period of 4 seconds (frequency 0.25 Hz). Methods to calculate seismic wavefields at such frequencies have continually evolved as the model used to represent Earth's crust and upper mantle has been refined. Current efforts to handle the complex scattering and wave propagation phenomena occurring between seismic source regions and Earth's surface rely primarily on finite difference or finite element computation. These methods are naturally able to account for general variations in elastic properties and are relatively simple to implement, but the amount of space demanded by them becomes prohibitive as the domain of wave propagation becomes larger.

We approach the problem of seismic scattering for strong ground motion applications using the analytic framework of coupled mode theory originally developed to perform very-low-frequency (0.01 – 0.05 Hz) seismic

wavefield computation. We have recently adapted these methods to higher frequency by introducing a refined technique for calculating the required wavefields, and are implementing them on LLNL supercomputers to simulate strong ground motion from crustal earthquakes, including the main effects of scattering through strongly heterogeneous crust and upper mantle structures. Our code can currently handle 3D variations in isotropic seismic wave velocities and density as well as attenuation. Previous research efforts in seismic scattering phenomena at lower frequency demonstrate the importance of including multifold mode-mode interactions in order to account for complex wave conversion effects over a broad range of spatial scales, as well as the importance of multiple scattering when large scattering domains are involved. Our work so far shows that these considerations are even more important at higher frequency. We intend to apply our new code to predict strong ground motion for various faulting scenarios in northern California, where combined body-wave and surface-wave based tomography now supply us with a satisfactory image of the highly variable crust and upper mantle structure.

Global Simulation of the Earth's Magnetosphere With Adaptive Mesh Refinement

Joachim Raeder

University of California, Los Angeles

Abstract

Our work adds structured adaptive mesh refinement capabilities to the UCLA Global Geospace Circulation Model (UCLA-GGCM) in collaboration with the SAMRAI project in the Center for Applied Scientific Computing at LLNL. The UCLA-GGCM is a global circulation model of Earth's magnetosphere and ionosphere that has been used for several years to study the interactions of the solar wind, magnetosphere, and ionosphere. The UCLA-GGCM with local adaptive mesh refinement supports a grid resolution in critical regions that is about two orders of magnitude better than the original code. This additional resolution allows us to address problems that were previously elusive to global modeling, such as flux transfer events, the formation of boundary layers, and the formation of thin current sheets in the late substorm growth phase.

SAMRAI is a general software support framework for structured adaptive mesh refinement (AMR) applications on parallel high performance computing hardware. To support local time refinement in the AMR time integration algorithm, we are investigating the numerical issues involved in local time synchronization of the magnetic and electric fluxes at coarse-fine grid interfaces.